Fair Petri Nets and Structural Induction for Rings of Processes*

Jianan Li
Department of Electrical Engineering and Computer Science
University of Wisconsin - Milwaukee
P.O. Box 784, Milwaukee, WI 53201, U.S.A.
jianan@convex.csd.uwm.edu

Ichiro Suzuki
Department of Electrical Engineering and Computer Science
University of Wisconsin - Milwaukee
P.O. Box 784, Milwaukee, WI 53201, U.S.A.
suzuki@cs.uwm.edu

Masafumi Yamashita
Department of Electrical Engineering
Faculty of Engineering
Hiroshima University
Kagamiyama, Higashi-Hiroshima 724, Japan
mak@se.hiroshima-u.ac.jp

To appear in *Theoretical Computer Science*
December 1994

Abstract We present a structural induction theorem for rings consisting of an arbitrary number of identical components. The components of a ring are modeled using a "fair Petri net," in which the firing of a prespecified set of transitions is assumed to occur fairly, i.e., any of these transitions that becomes firable infinitely often must fire infinitely often. Specifically, we introduce the concept of similarity between rings of different sizes, and give a condition under which the similarity between the rings of sizes two and three guarantees the similarity among the rings of all sizes. So if the given condition is satisfied, then the correctness of a ring of any large size can be inferred from the correctness of a ring having only a few components. The usefulness of the theorem is demonstrated using the examples of token-passing mutual exclusion and a simple producer-consumer system.

Send all correspondence to: Ichiro Suzuki, (414) 229-3718

*This work was supported in part by the National Science Foundation under grants CCR-9004346 and IRI-9307506, and the Office of Naval Research under grant N00014-94-1-0284. An extended abstract of an earlier version of this paper appeared as "Temporal Petri Nets and Structural Induction for Rings of Processes" in *Proceedings of the 35th Midwest Symposium on Circuits and Systems*, George Washington University, Washington, DC, August 9-12, 1992.

1 Introduction

Concurrent processing systems can exhibit extremely complicated behavior because of the complex timing of actions of different processes. Obtaining useful frameworks for analyzing such systems has been one of the major research problems in computer science.

In recent years, a number of papers have appeared that discuss the problem of analyzing concurrent systems consisting of a large number of finite state machines [1] [3] [5] [8] [20] [25]. The basic question there is to decide, given a system $S(n)$ consisting of $n \geq 2$ finite state machines and a property $P(n)$ on $S(n)$, whether or not $S(n)$ satisfies $P(n)$ for all values of n. Note that conventional theorem provers based on state-space search cannot be used directly to answer this question, since they can be applied only to instances having a fixed state-space. The impossibility of solving this problem in general was first shown by Apt and Kozen [1], and then Suzuki [20] sharpened the result by showing that the problem remains unsolvable even if $S(n)$ is a unidirectional ring of n identical finite state machines whose configuration is independent of the value of n. The results reported in [3] [5] [8] [25] are some of the efforts to find a sufficient condition for guaranteeing that $S(n)$ satisfies $P(n)$ for all values of n.

In this paper, we investigate the analysis problem stated above for systems that are *rings* of identical components, using *fair Petri nets* for representing the components. Intuitively, a fair Petri net is a Petri net in which the firing of a prespecified set of transitions is assumed to occur fairly, i.e., any of these transitions that becomes firable infinitely often must fire infinitely often. Formally, we define fair Petri nets as a subclass of *temporal Petri nets* [19]. Temporal Petri nets are Petri nets whose certain temporal constraints are given by formulas containing temporal operators, such as $\Diamond$ ("eventually") and $\Box$ ("always") [11] [12] [17]. Petri nets (see, for example, [14]) are widely used for modeling and analysis of concurrent processing systems. The combination of Petri nets and temporal logic has been found to be extremely useful for formal analysis of such systems [10] [21] [22] [23]. Theoretical studies of various temporal logic for Petri nets are found in [2] [6] [7] [19] [22] [23].

The main result of the paper is a structural induction theorem that can be used to formally infer the correctness of a ring of any large size from the correctness of a ring having only a few components. The theorem actually gives a sufficient condition for the "behavior" of a ring of any large size to be "similar" to that of a ring having only a few components. Specifically, for $k \geq 2$ let R^k be the ring consisting of k components. We define a concept of "similarity" for rings, and then show that if R^2 and R^3 are similar in this sense and certain additional conditions are satisfied, then for any $k \geq 4$, R^2 and R^k are also similar. This, together with the "correctness" of R^2 in a certain sense, can be used to ensure that R^k is also correct for all $k \geq 3$. Though the theorem is applicable only when R^k is bounded (i.e., the net representing R^k has only finitely many distinct reachable markings) for any $k \geq 2$, we give a weak sufficient condition for R^k to be bounded for any $k \geq 2$. (All the examples we discuss in the paper satisfy this condition.) The condition, which is given using the concept of an S-invariant [14], can be tested easily. In principle, if R^2 and R^3 are bounded then the similarity of R^2 and R^3 and the correctness of R^2 can be tested using an automatic theorem prover. The usefulness of the theorem is demonstrated using the well-known examples of token-passing mutual exclusion [16] and a simple producer-consumer system. Specifically, using the induction theorem we prove that the given algorithms for these problems guarantee certain liveness and safeness properties in R^k, regardless of the value of k.

The condition that a ring is bounded simply means that the ring is a finite state machine. Since all related papers mentioned above consider only systems consisting of finite state machines, the fact that our theorem can be applied only to bounded rings is not a severe restriction.

Our work has been inspired by those of Kurshan and McMillan [8] and Wolper and Lovinfosse [25] that present similar induction theorems. A common requirement of their induction methods is that the human verifier must first find an "invariant" (called "process invariant" or "network invariant") to carry out the induction. One difficulty in this approach is that finding such an invariant is not always easy (even if it exists). The method given in [3] that requires the establishment of a "bisimulation" between two systems seems to suffer from the same difficulty. In a sense, our induction theorem gives a sufficient condition for the existence of such an invariant (or bisimulation). Whether or not the condition of our theorem is satisfied can be tested using an automatic verifier (if the ring is bounded) and if so, the theorem assures the correctness of a ring of any size, given the correctness of a ring having a few components. There is no need for the human verifier to find an invariant to carry out the verification. It should also be mentioned, however, that the invariant method can be considered to be more general than ours, since it is possible that the condition of our theorem does not hold while a suitable invariant exists.

The rest of the paper is organized as follows. In Section 2 we review the basic terminology of Petri nets and temporal logic. The induction theorem is presented in Section 3 and then applied to the verification of two examples in Sections 4 and 5. The concluding remarks are found in Section 6.

2 Fair Petri Nets

The material presented in this section is basically the same as that given in Section 2 of [21].

For any set S, S^* is the set of all *finite* sequences of elements of S, including the empty sequence λ. S^ω denotes the set of all *infinite* sequences of elements of S. For a finite sequence $\alpha \in S^*$ and a possibly infinite sequence $\beta \in S^* \cup S^\omega$, $\alpha\beta$ denotes the concatenation of α and β. $\alpha\beta$ is an infinite sequence if β is an infinite sequence. $\alpha\beta$ is not defined if α is an infinite sequence. For a finite sequence $\alpha \in S^*$ and an integer $i \geq 0$, α^i denotes the concatenation of i copies of α. α^ω denotes the concatenation of infinitely many copies of α. $|\alpha|$ denotes the length of $\alpha \in S^*$. By convention we denote the length $|\alpha|$ of $\alpha \in S^\omega$ by ω, where ω is a symbol such that $i < \omega$ for any integer i.

A Petri net is a directed graph with two types of nodes, called transitions and places, and weighted arcs from a node of one type to a node of the other type. Formally, a *Petri net* is given as a triple $N = (P, T, F)$ where

1. P is a finite set of *places*,

2. T is a finite set of *transitions*, and

3. $F : (P \times T) \cup (T \times P) \rightarrow \{0, 1, 2, \ldots\}$ is a *weight* function.

A place $p \in P$ is called an *input place* (or *output place*) of a transition $t \in T$ if $F(p,t) \geq 1$ (or $F(t,p) \geq 1$). Any function $M : P \rightarrow \{0, 1, 2, \ldots\}$ is called a *marking*. A place p is said to have $M(p)$ [illegible]ens at a marking M. A transition $t \in T$ is said to be *firable* at M iff $M(p) \geq F(p$ [illegible] every $p \in P$. If t is firable at M, then it may *fire* and yield another marking M' [illegible] that $M'(p) = M(p) - F(p,t) + F(t,p)$ for every $p \in P$. We denote this by $M \rightarrow_t M'$. This relation is extended by

1. $M \rightarrow_\lambda M$ and

2. $M \rightarrow_{\alpha t} M'$ iff there exists M'' such that $M \rightarrow_\alpha M''$ and $M'' \rightarrow_t M'$

for all M, M', $\alpha \in T^*$ and $t \in T$. If $M \rightarrow_\alpha M'$ then M' is said to be *reachable* from M by a *finite firing sequence* α. $L(N,M)$ denotes the set of all finite firing sequences from M. An infinite sequence $\alpha \in T^\omega$ is an *infinite firing sequence* [24] from M if $\beta \in L(N,M)$ for every prefix β of α. We denote by $L^\omega(N,M)$ the set of infinite firing sequences from M. Let $L^\infty(N,M) = L(N,M$ [illegible] $,M)$ denote the set of *all* (both finite and infinite) firing sequences from M. Petr[illegible] *structurally bounded* if for any marking M, there are only finitely many distinct markings reachable from M. Usually an *initial marking* is associated with a Petri net.

We draw a Petri net using a circle and a square to represent places and transitions, respectively. An arc with weight $F(p,t)$ (or $F(t,p)$) is drawn from p to t (or from t to p) if $F(p,t) \geq 1$ (or $F(t,p) \geq 1$). The weight is omitted if it is 1. A marking M is represented by drawing $M(p)$ dots in (the circle representing) p. Examples of Petri nets are found in Section 3.

A *temporal Petri net* [19] [22] is a pair (N,f) where $N = (P,T,F)$ is a Petri net and f is a formula.[1] The formula f is regarded as a restriction on the possible firing sequences of N. For a marking M, we denote by $\mathcal{L}(N,M,f)$ the set of firing sequences $\alpha \in L^\infty(N,M)$ such that

1. α is either infinite, or finite and *terminating* in the sense that there is no transition $t \in T$ such that $\alpha t \in L(N,M)$, and

2. α satisfies f.

The first condition given above implies that the net is assumed to make progress whenever possible. In this paper we only consider formulas having the form

$$f(T') = \bigwedge_{t \in T'} ((\Box\Diamond \uparrow t) \supset (\Box\Diamond t)), \tag{1}$$

where $T' \subseteq T$ is a subset of transitions. We call such $f(T')$ an *f-formula*, where 'f' stands for "fairness," since an infinite sequence α satisfies $f(T')$ iff every $t \in T'$ that becomes firable infinitely often ($\Box\Diamond \uparrow t$) in α fires infinitely often ($\Box\Diamond t$) in α. The transitions in $T - T'$ need not be fired fairly. For example, if we wish to allow the system to issue a request for entering the critical section only a finite number of times, then the transition representing the action of making such a request may be excluded from T'. We call a temporal Petri net having a formula of the form (1) a *fair Petri net.*

Let $\overline{\mathcal{L}}(N,M,f)$ be the set of all prefixes of the sequences in $\mathcal{L}(N,M,f)$.

Lemma 1 *If f is an f-formula, then* $\overline{\mathcal{L}}(N,M,f) = L(N,M)$.

[1] See [19] [21] [22] for a formal discussion on the formulas.

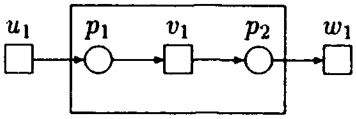

Figure 1: A component having one interface transition on each side.

Proof Clearly $\overline{\mathcal{L}}(N,M,f) \subseteq L(N,M)$. Since f is an f-formula, for any $\alpha \in L(N,M)$ there exists some β such that $\alpha\beta \in \mathcal{L}(N,M,f)$, and hence $\alpha \in \overline{\mathcal{L}}(N,M,f)$. Therefore $\overline{\mathcal{L}}(N,M,f) \supseteq L(N,M)$. □

3 Structural Induction on a Ring

In this section, we present a structural induction theorem that can be used to prove the correctness of rings of many similar components that are modeled as fair Petri nets. It is well-known that such induction is not always possible [1] [20]. The theorem presented here gives a sufficient condition under which the correctness of a ring of any large size can be inferred from the correctness of rings having only a few components.

Definition 1 A *component* is a Petri net $C = (P,T,F)$ in which the set T of transitions can be partitioned as $T = T_{\mathrm{L}} \cup T_{\mathrm{I}} \cup T_{\mathrm{R}}$ such that $|T_{\mathrm{L}}| = |T_{\mathrm{R}}| \geq 1$. The transitions in T_{L}, T_{I} and T_{R} are called *left interface transitions*, *internal transitions* and *right interface transitions*, respectively.

Figure 1 shows a component having one left interface transition u_1, one right interface transition w_1, one internal transition v_1, and two places p_1 and p_2.

We connect two or more components to form either a chain or a ring by merging the interface transitions of different components. The internal transitions of a component do not directly participate in the communication with other components. Formally, we have the following definitions.

Definition 2 Let $C = (P,T,F)$ be a component having places $P = \{p_1,\ldots,p_n\}$, left interface transitions $T_{\mathrm{L}} = \{u_1,\ldots,u_m\}$, internal transitions $T_{\mathrm{I}} = \{v_1,\ldots,v_s\}$, and right interface transitions $T_{\mathrm{R}} = \{w_1,\ldots,w_m\}$, where $T = T_{\mathrm{L}} \cup T_{\mathrm{I}} \cup T_{\mathrm{R}}$. For each $i \geq 0$, $C_i = (P_i, T_i, F_i)$ denotes the Petri net having the same structure as C in which

1. each p_j is renamed $p_{i,j}$, $P_i = \{p_{i,1},\ldots,p_{i,n}\}$,
2. each u_j is renamed $t_{i-1,j}$, $T_{i,\mathrm{L}} = \{t_{i-1,1},\ldots,t_{i-1,m}\}$,
3. each v_j is renamed $v_{i,j}$, $T_{i,\mathrm{I}} = \{v_{i,1},\ldots,v_{i,s}\}$,
4. each w_j is renamed $t_{i,j}$, $T_{i,\mathrm{R}} = \{t_{i,1},\ldots,t_{i,m}\}$,

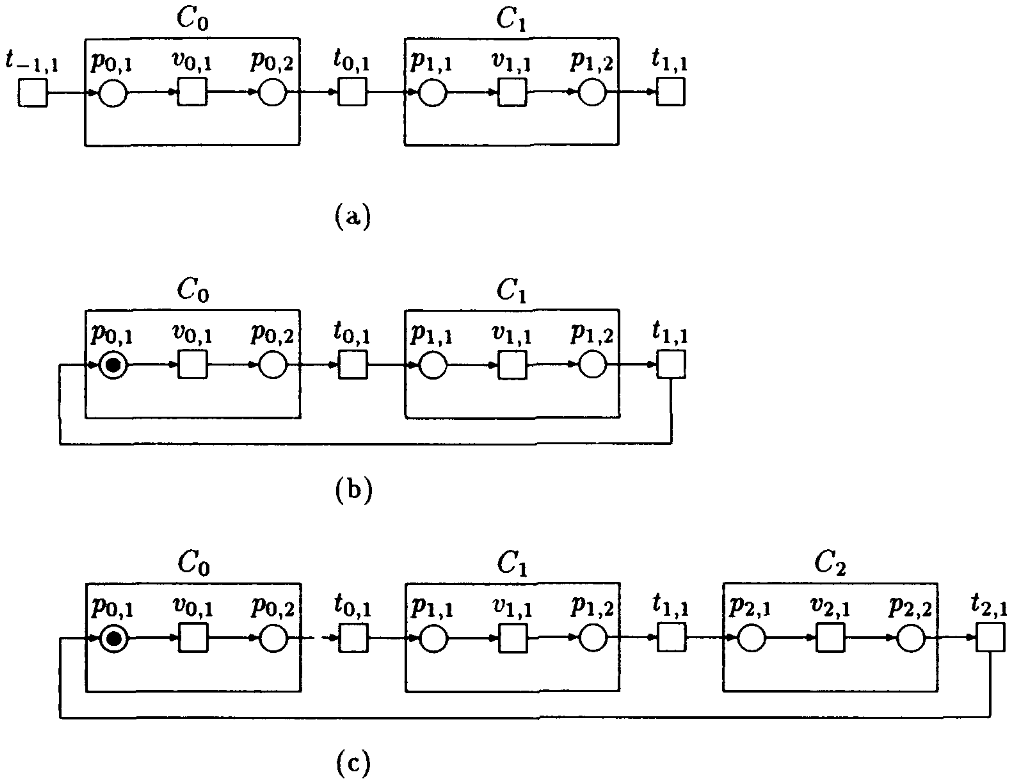

Figure 2: (a) $C_0 \oplus C_1$, (b) $R^2 = C_0 \odot C_1$ and (c) $R^3 = C_0 \odot C_1 \oplus C_2$ consisting of the component of Figure 1.

and $T_i = T_{i,\mathrm{L}} \cup T_{i,\mathrm{I}} \cup T_{i,\mathrm{R}}$. F_i is identical to F under the renaming given above. For $0 \le i \le j$,

$$C_i \oplus C_{i+1} \oplus \cdots \oplus C_j = \Big(\bigcup_{i \le \ell \le j} P_\ell,\ \bigcup_{i \le \ell \le j} T_\ell,\ \bigcup_{i \le \ell \le j} F_\ell \Big)$$

denotes the chain consisting of $C_i, C_{i+1}, \ldots, C_j$. (Note, for example, that C_i's right interface transitions, $t_{i,1}, \ldots, t_{i,m}$, have the same names as the left interface transitions of C_{i+1}. So in $C_i \oplus C_{i+1} \oplus \cdots \oplus C_j$, C_i and C_{i+1} are connected through $t_{i,1}, \ldots, t_{i,m}$.) For each $k \ge 2$,

$$R^k = C_0 \odot C_1 \oplus \cdots \oplus C_{k-1} = \Big(\bigcup_{0 \le \ell \le k-1} P_\ell,\ \bigcup_{0 \le \ell \le k-1} T_\ell,\ \bigcup_{0 \le \ell \le k-1} F_\ell \Big)$$

where all subscripts are taken modulo k, denotes the ring consisting of $C_0, C_1, \ldots, C_{k-1}$.

See Figure 2 for illustration. (Ignore the tokens at this time.) Chain $C_i \oplus C_{i+1} \oplus \cdots \oplus C_j$ is viewed as a new component having left interface transitions $T_{i,\mathrm{L}}$ and right interface transitions $T_{j,\mathrm{R}}$. The symbol "$\odot$" in $C_0 \odot C_1 \oplus \cdots \oplus C_{k-1}$ can be viewed as an operator that closes the chain $C_1 \oplus \cdots \oplus C_{k-1}$ into a ring using C_0, where $\oplus$ has precedence over $\odot$. All subscripts are taken modulo k when we discuss R^k. So for example, the left interface transitions of C_0 in R^k are $t_{k-1,1}, \ldots, t_{k-1,m}$, and $C_3 \oplus C_4 \oplus C_5 \oplus C_0 \oplus C_1$ is the chain embedded in R^6 consisting of C_3, C_4, C_5, C_0 and C_1. For each $0 \leq i \leq k-1$, we let $I_i = \{t_{i,1}, \ldots, t_{i,m}\}$ denote the set of interface transitions between C_i and C_{i+1}. The internal transitions of C_i and the interface transitions in $I_{i-1} \cup I_i$ are said to *belong to* C_i. An interface transition thus belongs to two components.

Since $C_0, C_1, \ldots, C_{k-1}$ are copies of C, a marking of R^k can simply be described as a tuple $(M_0, M_1, \ldots, M_{k-1})$, where each M_i is a marking of C, so that the number of tokens in $p_{i,j}$ of C_i is given by $M_i(p_j)$. We assume that all components of a ring except possibly C_0 have the same initial marking. As is the case with token-passing mutual exclusion [16], it is sometimes necessary that we break symmetry by assigning a different initial marking to one component. Thus for some markings M and M' of C, we let

$$M^k = (M, \underbrace{M', \ldots, M'}_{k-1})$$

be the initial marking of R^k in which C_0 has marking M and $C_1, \ldots, C_{k-1}$ have M'.

To ensure that the fairness requirement is imposed on an identical set of transitions at every component of a ring, we take T^k to be a set of transitions of R^k such that

1. for each $1 \leq j \leq s$, either $v_{i,j} \in T^k$ for all $0 \leq i \leq k-1$ or $v_{i,j} \notin T^k$ for all $0 \leq i \leq k-1$, and

2. for each $1 \leq j \leq m$, either $t_{i,j} \in T^k$ for all $0 \leq i \leq k-1$ or $t_{i,j} \notin T^k$ for all $0 \leq i \leq k-1$,

and then let $f^k = f(T^k)$ be an f-formula for R^k having the form (1). For a transition t, we say that $\alpha \in L^\infty(R^k, M^k)$ is *t-legal* if either it is infinite and satisfies $((\Box\Diamond \uparrow t) \supset (\Box\Diamond t))$, or it is finite and terminating. α is said to be *legal at* C_i if it is t-legal for all transitions $t \in T^k$ that belong to C_i. Note that α belongs to $\mathcal{L}(R^k, M^k, f^k)$ iff α is legal at every C_i.

In the following, sets $L(R^k, M^k)$, $L^\omega(R^k, M^k)$, $L^\infty(R^k, M^k)$ and $\mathcal{L}(R^k, M^k, f^k)$ are simply written as $L(k)$, $L^\omega(k)$, $L^\infty(k)$ and $\mathcal{L}(k)$, respectively. For convenience, we use "R^k" to refer to either the Petri net R^k alone or the tuple (R^k, M^k, f^k), depending on the context. There will be no confusion.

Remark 1 Since the initial marking M of C_0 can be different from those (M') of $C_1, \ldots, C_{k-1}$, C_0 can behave completely differently from $C_1, \ldots, C_{k-1}$. Thus many of the results presented below can be extended to the case when the structure and formula of C_0 are different from those of $C_1, \ldots, C_{k-1}$. In this paper, we assume that C_0 is the same as other components for simplicity of presentation.

Lemma 2 *If R^2 is structurally bounded, then for any $k \geq 3$, R^k is structurally bounded.*

Proof The proof is found in Appendix A. □

A place of a component that is an input place of a left (or right) interface transition is called a *left* (or *right*) *interface place* of the component. Since a chain $C_i \oplus \cdots \oplus C_j$ is viewed as a component, its left (or right) interface places are the left (or right) interface places of C_i (or C_j).

Definition 3 Let $p_{i_1}, \ldots, p_{i_L}$ and $p_{j_1}, \ldots, p_{j_R}$ be the left and right interface places of C, respectively, where $1 \leq i_1 < \cdots < i_L \leq n$ and $1 \leq j_1 < \cdots < j_R \leq n$. Then for a chain $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ of length at most $k-1$ within R^k and a marking $(M_0, M_1, \ldots, M_{k-1})$, the *firability vector* of $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ at $(M_0, M_1, \ldots, M_{k-1})$ is the column vector

$$\begin{bmatrix} M_{a-1}(p_{j_1}) \\ \vdots \\ M_{a-1}(p_{j_R}) \\ M_a(p_{i_1}) \\ \vdots \\ M_a(p_{i_L}) \\ M_b(p_{j_1}) \\ \vdots \\ M_b(p_{j_R}) \\ M_{b+1}(p_{i_1}) \\ \vdots \\ M_{b+1}(p_{i_L}) \end{bmatrix}.$$

Whether or not an interface transition (in $I_{a-1} \cup I_b$) of chain $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ is firable at marking $(M_0, M_1, \ldots, M_{k-1})$ can be determined by examining the firability vector of $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ at that marking.

Definition 4 Let $\alpha = t_1 t_2 \ldots t_i \ldots \in L^\infty(k)$ be a firing sequence such that for each $0 \leq i \leq |\alpha|$, $M^k \rightarrow_{t_1 t_2 \ldots t_i} M_i^k$. (Thus $M^k = M_0^k$.) For an index $0 \leq a \leq k-1$, let V_i be the firability vector of C_a at M_i^k. The *extended local history* of C_a in α, denoted $\langle C_a \rangle_\alpha$, is the sequence obtained from $V_0 t_1 V_1 t_2 V_2 \ldots$ by

1. deleting all transitions that do not belong to C_a,
2. replacing every remaining $t_{a-1,j}$, $v_{a,j}$ and $t_{a,j}$ by u_j, v_j and w_j, respectively, and then
3. replacing every maximal substring of identical vectors $V_{i_1} V_{i_2} \ldots V_{i_\ell}$ by a single occurrence of V_{i_1}.

The *local history* of C_a in α, denoted $\langle\langle C_a \rangle\rangle_\alpha$, is the firing sequence obtained from $\langle C_a \rangle_\alpha$ by deleting the firability vectors.

$\langle\langle C_a \rangle\rangle_\alpha$ is the firing sequence of C corresponding to the portion of α that occurs in C_a. $\langle C_a \rangle_\alpha$ is $\langle\langle C_a \rangle\rangle_\alpha$ together with the information on all the changes in the firability vector of C_a. We define

$$h(\alpha) = (\langle C_0 \rangle_\alpha, \langle C_1 \rangle_\alpha, \ldots, \langle C_{k-1} \rangle_\alpha).$$

In the following, if α is legal at C_a, then we say that $\langle C_a \rangle_\alpha$ (or $\langle\langle C_a \rangle\rangle_\alpha$) is legal.

Definition 5 For an index $0 \leq a \leq k-1$ and a firing sequence $\alpha \in L^\infty(k)$, the *externally visible history* of C_a in α, denoted $[C_a]_\alpha$, is the sequence obtained from $\langle C_a \rangle_\alpha$ by

1. deleting all the internal transitions of C, and then
2. replacing every maximal substring of identical vectors $V_{i_1}V_{i_2}\ldots V_{i_\ell}$, if any, by a single occurrence of V_{i_1}.

$[C_a]_\alpha$ is the firing sequence of the interface transitions of C corresponding to the portion of α that occurs in C_a, together with the information on all the changes in its firability vector. Since a chain is viewed as a component, we extend the concept of externally visible history of a single component to that of a chain. Thus for chain $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ and $\alpha \in L^\infty(k)$,

$$[C_a \oplus C_{a+1} \oplus \cdots \oplus C_b]_\alpha$$

is the sequence showing the firing of its interface transitions in $I_{a-1} \cup I_b$ and all the changes in its firability vector. (Note that we use $\langle\langle C_a\rangle\rangle_\alpha$ to denote the firing sequence of C corresponding to the transition firings in C_a in α, and thus $\langle\langle C_a \oplus C_{a+1} \oplus \cdots \oplus C_b\rangle\rangle_\alpha$ is not defined unless the chain consists of a single component. Similarly, $\langle C_a \oplus C_{a+1} \oplus \cdots \oplus C_b\rangle_\alpha$ is not defined unless the chain consists of a single component.)

Example 1 Consider rings R^k consisting of the component of Figure 1. Assume that at the initial marking M^k, place $p_{0,1}$ (the copy of p_1 in C_0) has one token and all other places are token-free. Figure 2 shows R^2 and R^3 with their initial markings. Since no two transitions share an input place, the fairness requirement is redundant. That is, we can take $T^k = \emptyset$ and $f^k = f(T^k) = \textbf{true}$. Then the only firing sequence in $\mathcal{L}(2)$ is

$$\alpha = (v_{0,1}t_{0,1}v_{1,1}t_{1,1})^\omega.$$

The firability vectors of C_0 have the form

$$\begin{bmatrix} x \\ y \end{bmatrix}$$

where x and y are the token counts of interface places $p_{1,2}$ and $p_{0,2}$, respectively. It is easy to show

$$\langle C_0\rangle_\alpha = (\begin{bmatrix} 0 \\ 0 \end{bmatrix} v_1 \begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1)^\omega,$$

$$\langle\langle C_0\rangle\rangle_\alpha = (v_1 w_1 u_1)^\omega$$

and

$$[C_0]_\alpha = (\begin{bmatrix} 0 \\ 0 \end{bmatrix} \begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1)^\omega.$$

In R^3, the firability vectors of C_0 show the token counts of interface places $p_{2,2}$ and $p_{0,2}$. The only firing sequence in $\mathcal{L}(3)$ is

$$\beta = (v_{0,1}t_{0,1}v_{1,1}t_{1,1}v_{2,1}t_{2,1})^\omega$$

and the reader can verify that $\langle C_0\rangle_\alpha = \langle C_0\rangle_\beta$, $\langle\langle C_0\rangle\rangle_\alpha = \langle\langle C_0\rangle\rangle_\beta$ and $[C_0]_\alpha = [C_0]_\beta$. □

Lemma 3 *Let $\alpha \in L^\infty(k)$ and $\beta \in L^\infty(\ell)$ be firing sequences such that for some $0 \leq a \leq b \leq k-1$ and $0 \leq c \leq d \leq \ell - 1$,*

$$[C_a \oplus C_{a+1} \oplus \cdots \oplus C_b]_\alpha = [C_c \oplus C_{c+1} \oplus \cdots \oplus C_d]_\beta.$$

Let $J = c + (b - a + 1) + (\ell - 1 - d) = \ell + (c - d) + (b - a)$. Then there exists a firing sequence $\gamma \in L^\infty(J)$ such that

$$h(\gamma) = (\underbrace{\langle C_0\rangle_\beta, \ldots, \langle C_{c-1}\rangle_\beta}_{c}, \underbrace{\langle C_a\rangle_\alpha, \ldots, \langle C_b\rangle_\alpha}_{b-a+1}, \underbrace{\langle C_{d+1}\rangle_\beta, \ldots, \langle C_{\ell-1}\rangle_\beta}_{\ell-1-d}).$$

Proof We only give an outline. Suppose that we construct a ring of size J by connecting $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ of R^k and $C_{d+1} \oplus \cdots \oplus C_{\ell-1} \oplus C_0 \oplus \cdots C_{c-1}$ of R^ℓ. Since $[C_a \oplus C_{a+1} \oplus \cdots \oplus C_b]_\alpha = [C_c \oplus C_{c+1} \oplus \cdots \oplus C_d]_\beta$, we can fire the transitions in α that belong to $C_a \oplus C_{a+1} \oplus \cdots \oplus C_b$ and the transitions in β that belong to $C_{d+1} \oplus \cdots \oplus C_{\ell-1} \oplus C_0 \oplus \cdots \oplus C_{c-1}$ in such a way that (a) the interface transitions between C_{c-1} and C_a, and between C_b and C_{d+1}, are fired simultaneously, and (b) the token counts of the input places of the interface transitions between C_{c-1} and C_a, and between C_b and C_{d+1}, change in the same manner as those of the input places of the interface transitions in $I_{a-1} \cup I_b$ in α. The resulting sequence γ is a firing sequence in $L^\infty(J)$ satisfying the condition on $h(\gamma)$ given above. □

Recall that $\mathcal{L}(k)$ is the set of firing sequences α in R^k from M^k satisfying f^k, i.e., α is legal at every C_i. For each $0 \leq i \leq k-1$, we denote by $\mathcal{L}_{\neg i}(k)$ the set of firing sequences $\alpha \in L^\infty(k)$ that are t-legal for all transitions t of R^k except possibly the internal transitions of C_i. Such α may or may not be legal at C_i.

Definition 6 $R^k = C_0 \odot C_1 \oplus \cdots \oplus C_{k-1}$ and $R^\ell = C_0 \odot C_1 \oplus \cdots \oplus C_{\ell-1}$ are *similar*, denoted $R^k \sim R^\ell$, if

1. $\{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ and
2. $\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}(\ell)\}$ for any $1 \leq i \leq k-1$ and $1 \leq j \leq \ell - 1$.

Definition 7 $R^k = C_0 \odot C_1 \oplus \cdots \oplus C_{k-1}$ and $R^\ell = C_0 \odot C_1 \oplus \cdots \oplus C_{\ell-1}$ are *strongly similar*, denoted $R^k \approx R^\ell$, if

1. $\{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}_{\neg 0}(k)\} = \{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}_{\neg 0}(\ell)\}$ and
2. $\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}_{\neg i}(k)\} = \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}_{\neg j}(\ell)\}$ for any $1 \leq i \leq k-1$ and $1 \leq j \leq \ell - 1$.

Intuitively, if $R^k \sim R^\ell$, then as long as the components behave legally, none of the copies of C knows which of R^k and R^ℓ it is in, and none of the copies of C other than C_0 knows which copy of C it is. The strong similarity $R^k \approx R^\ell$ assures that the same is true for any copy of C that may violate the f-formula for its internal transitions, as long as all other components behave legally. Note that $R^k \approx R^\ell$ implies $R^k \sim R^\ell$.

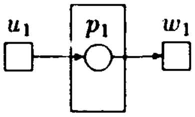

Figure 3: A component having one place.

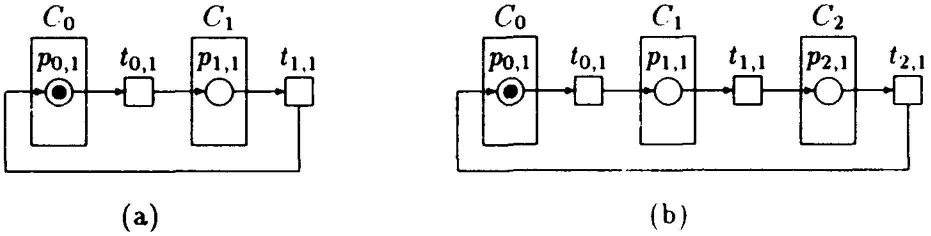

Figure 4: (a) R^2 and (b) R^3 consisting of the component of Figure 3.

Example 2 We have seen that rings R^2 and R^3 given in Example 1 satisfy

$$\{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(2)\} = \{\langle C_0\rangle_\alpha | \alpha \in \mathcal{L}(3)\}.$$

Using a similar argument, we can also show that

$$\{\langle C_1\rangle_\alpha | \alpha \in \mathcal{L}(2)\} = \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}(3)\}$$

for $j = 1, 2$. Thus $R^2 \sim R^3$. We leave it to the reader to verify that in fact, $R^2 \sim R^k$ holds for any $k \geq 3$. Furthermore, since $T^k = \emptyset$, $\mathcal{L}_{\neg i}(k) = \mathcal{L}(k)$ for any $0 \leq i \leq k-1$. Therefore $R^2 \sim R^k$ implies $R^2 \approx R^k$. □

Example 3 Consider rings R^2 and R^3 shown in Figure 4 consisting of the component of Figure 3. Assume that at the initial marking, place $p_{0,1}$ has one token and all other places are token-free. As in Example 1, take $T^k = \emptyset$, and thus $f^k = f(T^k) = \textbf{true}$. The only firing sequence in $\mathcal{L}(2)$ is

$$\alpha = (t_{0,1}t_{1,1})^\omega$$

with

$$\langle C_0\rangle_\alpha = (\begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1 \begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1)^\omega,$$

where the firability vectors of C_0 show the token counts of interface places $p_{1,1}$ and $p_{0,1}$. As for R^3, the only firing sequence in $\mathcal{L}(3)$ is

$$\alpha = (t_{0,1}t_{1,1}t_{2,1})^\omega$$

with

$$\langle C_0\rangle_\beta = (\begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1 \begin{bmatrix} 0 \\ 0 \end{bmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1)^\omega,$$

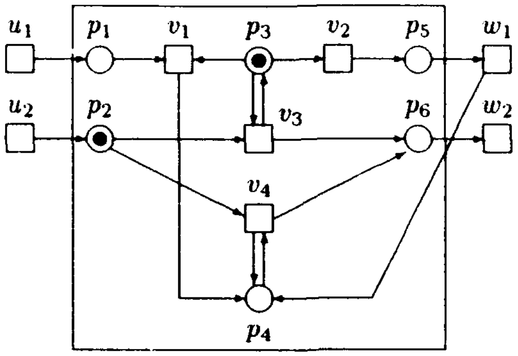

Figure 5: Component C such that $R^2 \sim R^3$ but $R^2 \not\approx R^3$.

where the firability vectors of C_0 show the token counts of interface places $p_{2,1}$ and $p_{0,1}$. Since $\langle C_0 \rangle_\alpha \neq \langle C_0 \rangle_\beta$, we have $R^2 \not\sim R^3$. On the other hand, it is easy to show that $R^3 \sim R^k$ holds for any $k \geq 4$. Since $T^k = \emptyset$, this implies that $R^3 \approx R^k$ for any $k \geq 4$. We leave details to the reader. □

Example 4 Figure 5 shows a component C such that $R^2 \sim R^3$ but $R^2 \not\approx R^3$. T^k is the set of all transitions in R^k. Initially, C_0 has a token in p_2 and p_3. (Strictly, we should say that C_0 has a token in $p_{0,2}$ and $p_{0,3}$, that are the copies of p_2 and p_3 in C_0. For convenience, in this example we use the original names in C to refer to places and transitions of C_i.) All other components C_i have a token only in p_3. Intuitively, the components keep circulating the token that is initially in p_2 of C_0, using u_2, v_3 and w_2, and later using u_2, v_4 and w_2 since $v_2 w_1$ should eventually fire to satisfy the fairness condition, unless $u_1 v_1$ fires. Suppose that in R^3, C_0 violates fairness and fires $v_3 w_2 (u_2 v_3 w_2)^\omega$. C_1 and C_2 can still continue to circulate the token indefinitely without violating fairness, by firing $v_2 w_1$ in C_1 and $u_1 v_1$ in C_2 and thus moving the token in p_3 to p_4 in both components. (Note that w_1 of C_1 is the same as u_1 of C_2.) In R^2, however, if C_0 violates fairness and fires $v_3 w_2 (u_2 v_3 w_2)^\omega$, then C_1 eventually fires v_2 (to satisfy fairness) but it cannot fire w_1, since w_1 of C_1 is the same as u_1 of C_0 and C_0 never fires u_1. So w_1 of C_1 remains firable forever and never fires, and thus fairness is violated at C_1. A formal analysis based on this observation shows that $R^2 \not\approx R^3$. The fact that such a scenario cannot happen if all components behave fairly is the basis for proving $R^2 \sim R^3$. We leave details to the reader. □

The main goal of this section is to prove the next theorem that can be used to prove the correctness of rings consisting of an arbitrary number of copies of C.

Theorem 1 *If R^2 is structurally bounded and $R^2 \approx R^3$, then $R^2 \sim R^k$ for any $k \geq 3$.*

We need the following lemmas to prove this theorem.

Lemma 4 *If $R^2 \sim R^3$ and $R^2 \sim R^k$ for some $k \geq 3$, then whenever either $i = j = 0$ or both $1 \leq i \leq k-1$ and $1 \leq j \leq k$,*

$$\{\langle C_i \rangle_\alpha | \alpha \in \mathcal{L}(k)\} \subseteq \{\langle C_j \rangle_\alpha | \alpha \in \mathcal{L}(k+1)\}.$$

Proof Since $R^2 \sim R^k$ implies that the sets $\{\langle C_i \rangle_\alpha | \alpha \in \mathcal{L}(k)\}$ are all identical for $1 \leq i \leq k-1$, it suffices to show that for any $\alpha \in \mathcal{L}(k)$, where

$$h(\alpha) = (\langle C_0 \rangle_\alpha, \langle C_1 \rangle_\alpha, \ldots, \langle C_{k-1} \rangle_\alpha),$$

there exist β and $\beta' \in \mathcal{L}(k+1)$ such that

$$h(\beta) = (\langle C_0 \rangle_\alpha, \langle C_1 \rangle_\alpha, \ldots, \langle C_{k-2} \rangle_\alpha, \underbrace{\langle C_{k-1} \rangle_\beta, \langle C_k \rangle_\beta}_{[C_{k-1}]_\alpha})$$ [2]

and

$$h(\beta') = (\underbrace{\langle C_0 \rangle_{\beta'}, \langle C_1 \rangle_{\beta'}}_{[C_0]_\alpha}, \langle C_1 \rangle_\alpha, \langle C_2 \rangle_\alpha, \ldots, \langle C_{k-1} \rangle_\alpha).$$

In the following we show the existence of such β. The argument for β' is similar and is thus omitted. Since $R^2 \sim R^k$, there exists $\gamma \in \mathcal{L}(2)$ such that

$$h(\gamma) = (\quad \underbrace{\langle C_0 \rangle_\gamma}_{[C_0 \oplus C_1 \oplus \cdots \oplus C_{k-2}]_\alpha} \quad, \langle C_{k-1} \rangle_\alpha).$$

Since $R^2 \sim R^3$, there exists $\delta \in \mathcal{L}(3)$ such that

$$h(\delta) = (\langle C_0 \rangle_\gamma, \underbrace{\langle C_1 \rangle_\delta, \langle C_2 \rangle_\delta}_{[C_{k-1}]_\alpha}).$$

Since $[C_{k-1}]_\alpha = [C_1 \oplus C_2]_\delta$, by Lemma 3 there exists $\epsilon \in L^\infty(k+1)$ such that

$$h(\epsilon) = (\langle C_0 \rangle_\alpha, \langle C_1 \rangle_\alpha, \ldots, \langle C_{k-2} \rangle_\alpha, \underbrace{\langle C_1 \rangle_\delta, \langle C_2 \rangle_\delta}_{[C_{k-1}]_\alpha}).$$

Since all elements of $h(\epsilon)$ are legal, ϵ satisfies f^{k+1}. Therefore $\epsilon \in \mathcal{L}(k+1)$. □

Remark 2 The proofs of Lemmas 3 and 4 do not use the assumption that f is an f-formula. In fact, the two lemmas are true for an arbitrary formula f, as long as the legality of any α is determined only by the legality of the elements of $h(\alpha)$.

Lemma 5 *Let t be a left (or right) interface transition of C_i of R^2 or R^3. If $R^2 \sim R^3$, then a firing of t does not change the token counts of the right (or left) interface places of C_i.*

Proof We consider the case when t is a left interface transition of C_1 of R^3. Other cases are similar. Take any $\alpha t \in \mathcal{L}(3)$. Since $R^2 \sim R^3$, there exists $\beta t \in \mathcal{L}(2)$ such that $\langle C_0 \rangle_{\alpha t} = \langle C_0 \rangle_{\beta t}$. Suppose that the firing of t in αt changes the token counts of the right interface places of C_1 of R^3. Then the firing of t in βt changes the token counts of the right interface places of C_1 of R^2, since C_1 has the same structure in R^2 and R^3. Then, since $\langle C_0 \rangle_{\alpha t} = \langle C_0 \rangle_{\beta t}$ implies $\langle C_0 \rangle_\alpha = \langle C_0 \rangle_\beta$, the firing of t in αt should also change the token counts of the right interface places of C_2 of R^3. But this is impossible, since t does not belong to C_2. □

[2] The underbrace indicates that $[C_{k-1} \oplus C_k]_\beta = [C_{k-1}]_\alpha$. Although this relation is implied by the forms of $h(\alpha)$ and $h(\beta)$, we use this notation to improve readability.

Lemma 6 *If $R^2 \sim R^3$ and $R^2 \sim R^k$ for some $k \geq 3$, then for any $\alpha \in L(k+1)$, where*

$$h(\alpha) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \langle C_{k-1}\rangle_\alpha, \langle C_k\rangle_\alpha),$$

there exists $\beta \in L(k)$ such that

$$h(\beta) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \underbrace{\langle C_{k-1}\rangle_\beta}_{[C_{k-1}\oplus C_k]_\alpha}). \tag{2}$$

Proof The proof is by induction. If $\alpha = \lambda$, then $\beta = \lambda \in L(k)$ satisfies equation (2) since both in R^k and R^{k+1}, initially C_0 has marking M and all other copies of C have M'. Assume that for some $\alpha \in L(k+1)$ such that $\alpha t \in L(k+1)$, there exists $\beta \in L(k)$ satisfying equation (2). There are three cases.

Case 1: t is an internal transition of $C_0 \oplus C_1 \oplus \cdots \oplus C_{k-2}$.

Clearly t is firable in R^k after β, i.e., $\beta t \in L(k)$. Since a firing of t can change the token counts of the places in $C_0 \oplus C_1 \oplus \cdots \oplus C_{k-2}$ only, and the changes are identical in R^k and R^{k+1}, βt has the property

$$h(\beta t) = (\langle C_0\rangle_{\alpha t}, \langle C_1\rangle_{\alpha t}, \ldots, \langle C_{k-2}\rangle_{\alpha t}, \underbrace{\langle C_{k-1}\rangle_{\beta t}}_{[C_{k-1}\oplus C_k]_{\alpha t}}).$$

Case 2: $t \in I_{k-2} \cup I_k$.

Consider the case when $t \in I_{k-2}$, i.e., $t = t_{k-2,j}$ for some $1 \leq j \leq m$. Since $[C_{k-1}]_\beta = [C_{k-1} \oplus C_k]_\alpha$, we have $\beta t \in L(k)$. Now we prove that $[C_{k-1}]_{\beta t} = [C_{k-1} \oplus C_k]_{\alpha t}$. A firing of t changes the token counts of the right (or left) interface places of C_{k-2} (or C_{k-1}) in the same way in R^k and R^{k+1}, and it does not change the token counts of the left interface places of C_0 in either ring since t does not belong to C_0. Also, it does not change the token counts of the right interface places of C_k of R^{k+1} since t does not belong to C_k. It remains to be shown that the token counts of the right interface places of C_{k-1} of R^k do not change by a firing of t. This follows from Lemma 5 and the fact that, by Lemma 1 and $R^2 \sim R^k$, there exists $\gamma t' \in L(2)$ such that $\langle C_{k-1}\rangle_{\beta t} = \langle C_1\rangle_{\gamma t'}$, where $t' = t_{0,j}$ is the interface transition of C_1 in R^2 corresponding to t of C_k in R^k. Therefore

$$h(\beta t) = (\langle C_0\rangle_{\alpha t}, \langle C_1\rangle_{\alpha t}, \ldots, \langle C_{k-2}\rangle_{\alpha t}, \underbrace{\langle C_{k-1}\rangle_{\beta t}}_{[C_{k-1}\oplus C_k]_{\alpha t}}).$$

The argument for the case when $t \in I_k$ is similar.

Case 3: t is an internal transition of $C_{k-1} \oplus C_k$.

By Lemma 1 and the assumption that $R^2 \sim R^k$, there exists $\gamma \in L(2)$ such that

$$h(\gamma) = (\langle C_0\rangle_\gamma, \langle C_{k-1}\rangle_\beta).$$

Then since $[C_{k-1}]_\beta = [C_{k-1} \oplus C_k]_\alpha$, by Lemma 3 there exists $\delta \in L(3)$ such that

$$\begin{aligned} h(\delta) &= (\langle C_0\rangle_\delta, \langle C_{k-1}\rangle_\delta, \langle C_k\rangle_\delta) \\ &= (\langle C_0\rangle_\gamma, \langle C_{k-1}\rangle_\alpha, \langle C_k\rangle_\alpha). \end{aligned}$$

Let t' be the internal transition of $C_1 \oplus C_2$ in R^3 corresponding to t, i.e., either $t = t_{k-1,j}$ and $t' = t_{1,j}$ for some $1 \leq j \leq m$, or $t = v_{k-1,j}$ and $t' = v_{1,j}$ or $t = v_{k,j}$ and $t' = v_{2,j}$ for some $1 \leq j \leq s$. Then since $\alpha t \in L(k+1)$, we have $\delta t' \in L(3)$ where

$$h(\delta t') = (\langle C_0\rangle_{\delta t'}, \langle C_{k-1}\rangle_{\alpha t}, \langle C_k\rangle_{\alpha t}).$$

Since $R^2 \sim R^3$, there exists $\epsilon \in L(2)$ such that

$$h(\epsilon) = (\langle C_0\rangle_{\delta t'}, \underbrace{\langle C_1\rangle_\epsilon}_{[C_{k-1}\oplus C_k]_{\alpha t}}).$$

Since $[C_1]_\epsilon = [C_{k-1} \oplus C_k]_{\alpha t}$, by Lemma 3 there exists $\zeta \in L(k)$ such that

$$h(\zeta) = (\langle C_0\rangle_{\alpha t}, \langle C_1\rangle_{\alpha t}, \ldots, \langle C_{k-2}\rangle_{\alpha t}, \underbrace{\langle C_1\rangle_\epsilon}_{[C_{k-1}\oplus C_k]_{\alpha t}}).$$

This completes the induction. □

Lemma 7 *If R^2 is structurally bounded, $R^2 \approx R^3$ and $R^2 \sim R^k$ for some $k \geq 3$, then whenever either $i = j = 0$ or both $1 \leq i \leq k-1$ and $1 \leq j \leq k$,*

$$\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}(k)\} \supseteq \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}(k+1)\}.$$

Proof Since $R^2 \sim R^k$ implies that the sets $\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}(k)\}$ are all identical for $1 \leq i \leq k-1$, it suffices to show that for any $\alpha \in \mathcal{L}(k+1)$, where

$$h(\alpha) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \langle C_2\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \langle C_{k-1}\rangle_\alpha, \langle C_k\rangle_\alpha),$$

there exist β and $\beta' \in \mathcal{L}(k)$ such that

$$h(\beta) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \underbrace{\langle C_{k-1}\rangle_\beta}_{[C_{k-1}\oplus C_k]_\alpha}) \tag{3}$$

and

$$h(\beta') = (\underbrace{\langle C_0\rangle_{\beta'}}_{[C_0\oplus C_1]_\alpha}, \langle C_2\rangle_\alpha, \ldots, \langle C_{k-1}\rangle_\alpha, \langle C_k\rangle_\alpha).$$

In the following we show the existence of such β. The argument for β' is similar and is thus omitted.

First, we show that there exists $\gamma \in L^\infty(k)$ such that

$$h(\gamma) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \underbrace{\langle C_{k-1}\rangle_\gamma}_{[C_{k-1}\oplus C_k]_\alpha}). \tag{4}$$

Sequence γ is just like β of equation (3), except that it may not be legal at C_{k-1}. There are two cases.

Case 1: In α, (a) the interface transitions in $I_{k-2} \cup I_k$ fire only a finite number of times and (b) the token counts of the interface places of $C_{k-1} \oplus C_k$ change only a finite number of times.

Let $\alpha_1 \in L(k+1)$ be the shortest sequence such that (a) $\alpha = \alpha_1\alpha_2$, (b) the interface transitions in $I_{k-2} \cup I_k$ do not fire in α_2, and (c) the token counts of the interface places of $C_{k-1} \oplus C_k$ do not change in α_2. Let $\alpha' = \alpha_1\alpha_2' \in L^\infty(k+1)$ be the sequence that is identical to α except that no transition in $C_{k-1} \oplus C_k$ fires after α_1. By Lemma 6, there exists $\beta_1 \in L(k)$ such that

$$h(\beta_1) = (\langle C_0\rangle_{\alpha_1}, \langle C_1\rangle_{\alpha_1}, \ldots, \langle C_{k-2}\rangle_{\alpha_1}, \underbrace{\langle C_{k-1}\rangle_{\beta_1}}_{[C_{k-1} \oplus C_k]_{\alpha_1}}).$$

By the assumption on α given above, we can extend β_1 to $\gamma = \beta_1\alpha_2' \in L^\infty(k)$, which clearly satisfies equation (4).

Case 2: In α, either (a) the interface transitions in $I_{k-2} \cup I_k$ fire infinitely often, or (b) the token counts of the interface places of $C_{k-1} \oplus C_k$ change infinitely many times.

Such α can be written as $\alpha = \sigma_1 x_1 \sigma_2 x_2 \ldots$, where $x_1, x_2, \ldots$ are the interface transitions of $C_{k-1} \oplus C_k$ and the transitions in $C_{k-1} \oplus C_k$ whose firings change the token counts of the interface places of $C_{k-1} \oplus C_k$. For each $\ell \geq 1$, let $\alpha_\ell = \sigma_1 x_1 \sigma_2 x_2 \ldots \sigma_\ell x_\ell$ be the prefix of α ending with x_ℓ. By Lemma 6, for each ℓ there exists $\beta_\ell \in L(k)$ such that

$$h(\beta_\ell) = (\langle C_0\rangle_{\alpha_\ell}, \langle C_1\rangle_{\alpha_\ell}, \ldots, \langle C_{k-2}\rangle_{\alpha_\ell}, \underbrace{\langle C_{k-1}\rangle_{\beta_\ell}}_{[C_{k-1} \oplus C_k]_{\alpha_\ell}}).$$

Since $[C_{k-1} \oplus C_k]_{\alpha_\ell} = [C_{k-1}]_{\beta_\ell}$, β_ℓ can be written as $\beta_\ell = \tau_1 y_1 \tau_2 y_2 \ldots \tau_\ell y_\ell$, where if x_i is an interface transition of $C_{k-1} \oplus C_k$ then y_i is the corresponding interface transition of C_{k-1} of R^k, and otherwise x_i and y_i respectively change the token counts of the interface places of $C_{k-1} \oplus C_k$ and C_{k-1} in the same way. Clearly, we may assume that for each $i \geq 1$, the internal transitions of $C_0 \oplus \cdots \oplus C_{k-2}$ fire in exactly the same way in σ_i and τ_i. Let M_i be the marking of R^{k+1} reached right after the firing of x_i, and N_i the marking of R^k reached right after the firing of y_i. Call a tuple of the form

$$(x_i, M_i, y_i, N_i; x_{i+1}, M_{i+1})$$

a *pattern*. Note that by Lemma 2 and the assumption that R^2 is structurally bounded, both R^k and R^{k+1} have only finitely many distinct reachable markings. Thus if ℓ is sufficiently large, then all patterns that appear in $\alpha_{\ell'}$ and $\beta_{\ell'}$ for any $\ell' > \ell$ appear in α_ℓ and β_ℓ. Then for $\alpha_{\ell+1} = \sigma_1 x_1 \ldots \sigma_\ell x_\ell \sigma_{\ell+1} x_{\ell+1}$, there exists a pattern

$$(x_j, M_j, y_j, N_j; x_{j+1}, M_{j+1}) = (x_\ell, M_\ell, y_\ell, N_\ell; x_{\ell+1}, M_{\ell+1}),$$

$j < \ell$, that appears in α_ℓ and β_ℓ. This means that in R^k, we can fire $\tau_{j+1} y_{j+1}$ after β_ℓ and obtain a sequence $\beta' = \beta_\ell \tau_{j+1} y_{j+1} \in L(k)$. Clearly

$$h(\beta') = (\langle C_0\rangle_{\alpha_{\ell+1}}, \langle C_1\rangle_{\alpha_{\ell+1}}, \ldots, \langle C_{k-2}\rangle_{\alpha_{\ell+1}}, \underbrace{\langle C_{k-1}\rangle_{\beta'}}_{[C_{k-1} \oplus C_k]_{\alpha_{\ell+1}}})$$

and R^k is at marking N_{j+1} after β'. Thus we can continue to extend β' in a similar manner and obtain $\gamma \in L^\omega(k)$ such that

$$h(\gamma) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \underbrace{\langle C_{k-1}\rangle_\gamma}_{[C_{k-1}\oplus C_k]_\alpha}).$$

This completes the proof of the existence of γ satisfying equation (4).

Since γ obtained above may not be legal at C_{k-1}, we convert it into a legal sequence $\beta \in \mathcal{L}(k)$. By $R^2 \sim R^k$ and Lemma 1, for each prefix γ' of γ there exists $\gamma'' \in L(2)$ such that

$$h(\gamma'') = (\underbrace{\langle C_0\rangle_{\gamma''}}_{[C_0\oplus\cdots\oplus C_{k-2}]_{\gamma'}}, \langle C_{k-1}\rangle_{\gamma'}).$$

Thus by Lemma 2, the structural boundedness of R^2, and an argument similar to the one given above for α, we can show that there exists $\delta \in L^\infty(2)$ such that

$$h(\delta) = (\underbrace{\langle C_0\rangle_\delta}_{[C_0\oplus\cdots\oplus C_{k-2}]_\gamma}, \langle C_{k-1}\rangle_\gamma).$$

Then, since $[C_{k-1}]_\gamma = [C_{k-1} \oplus C_k]_\alpha$, by Lemma 3 there exists $\epsilon \in L^\infty(3)$ such that

$$h(\epsilon) = (\langle C_0\rangle_\delta, \langle C_{k-1}\rangle_\alpha, \langle C_k\rangle_\alpha).$$

Then since $R^2 \approx R^3$ and both $\langle C_{k-1}\rangle_\alpha$ and $\langle C_k\rangle_\alpha$ are legal, there exists $\zeta \in L^\infty(2)$ such that

$$h(\zeta) = (\langle C_0\rangle_\delta, \underbrace{\langle C_{k-1}\rangle_\zeta}_{[C_{k-1}\oplus C_k]_\alpha})$$

where $\langle C_{k-1}\rangle_\zeta$ is legal. Then by Lemma 3, there exists $\beta \in L^\infty(k)$ such that

$$h(\beta) = (\langle C_0\rangle_\alpha, \langle C_1\rangle_\alpha, \ldots, \langle C_{k-2}\rangle_\alpha, \underbrace{\langle C_{k-1}\rangle_\zeta}_{[C_{k-1}\oplus C_k]_\alpha}).$$

Since all elements of $h(\beta)$ including $\langle C_{k-1}\rangle_\zeta$ are legal, $\beta \in \mathcal{L}(k)$. This completes the proof of the existence of $\beta \in \mathcal{L}(k)$ satisfying equation (3). □

Proof of Theorem 1 By Lemmas 4 and 7, if R^2 is structurally bounded, $R^2 \approx R^3$ and $R^2 \sim R^k$ for some $k \geq 3$, then $R^2 \sim R^{k+1}$. Thus the theorem follows by induction. □

A typical argument for proving the correctness of R^k is to show that

$$S \subseteq \{\langle\langle C_i\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} \subseteq S'$$

holds for all $0 \leq i \leq k-1$, where S and S' are sets of firing sequences of C describing certain properties of C_i. (We may have to use slightly different sets for $i = 0$, since the initial marking of C_0 can be different from those of the other copies of C.) For example, S' may consist of the sequences in which every firing of a transition representing "request critical section" is followed by a firing of another transition representing "enter critical section," to ensure that every request of C_i to enter its critical section will eventually be granted. The use of some nonempty S eliminates the case when C_i satisfies the condition imposed by S' by having, for example, $\{\langle\langle C_i\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\} = \emptyset$. If R^2 is structurally bounded, $R^2 \approx R^3$ and R^2 is correct in the above sense, then by Theorem 1 and the fact that $\{\langle C_i\rangle_\alpha | \alpha \in \mathcal{L}(2)\} = \{\langle C_j\rangle_\alpha | \alpha \in \mathcal{L}(k)\}$ implies $\{\langle\langle C_i\rangle\rangle_\alpha | \alpha \in \mathcal{L}(2)\} = \{\langle\langle C_j\rangle\rangle_\alpha | \alpha \in \mathcal{L}(k)\}$, we can conclude that R^k is correct for any $k \geq 2$. In principle, if R^2 and R^3 are finite state systems, then the correctness of R^2 and whether or not $R^2 \approx R^3$ can be tested automatically using a conventional theorem prover.[3] As is seen from the discussion given in Appendix A, whether or not R^2 is structurally bounded can be tested by solving a set of linear inequalities.

Note that the proof method described above allows us to verify only "local" properties of the copies of C in R^k. To prove certain "global" properties of R^k, such as mutual exclusion ("only one copy of C in R^k can enter a critical section at a time"), we need a result such as the following.

For a firing sequence $\alpha \in \mathcal{L}(k)$ and each $0 \leq j \leq m$, where m is the number of interface transitions of C on each side, let $\rho_j(\alpha)$ be the sequence obtained from α by deleting all transitions except the j-th interface transitions $t_{0,j}, t_{1,j}, \ldots, t_{k-1,j}$.

Theorem 2 *If $R^k \sim R^{k+1}$ for some $k \geq 2$ and $\rho_j(\alpha)$ is either $(t_{0,j}t_{1,j}\ldots t_{k-1,j})^\omega$ or its prefix for any $\alpha \in \mathcal{L}(k)$, then $\rho_j(\alpha)$ is either $(t_{0,j}t_{1,j}\ldots t_{k-1,j}t_{k,j})^\omega$ or its prefix for any $\alpha \in \mathcal{L}(k+1)$.*

Proof Suppose that there exists $\alpha \in \mathcal{L}(k+1)$ such that $\rho_j(\alpha)$ is not $(t_{0,j}t_{1,j}\ldots t_{k-1,j}t_{k,j})^\omega$ or its prefix. Then in α, either (a) some component C_i, $i \neq 0$, fires $t_{i,j}$ before $t_{i-1,j}$ fires for the first time, or (b) some component C_i fires $t_{i,j}$ twice without firing $t_{i-1,j}$ between the two firings of $t_{i,j}$. (Here, subscript $i-1$ is computed mod $(k+1)$.) Since $R^k \sim R^{k+1}$, there exists $\beta \in \mathcal{L}(k)$ such that (a) $\langle C_i\rangle_\alpha = \langle C_1\rangle_\beta$ if $i \neq 0$ and (b) $\langle C_i\rangle_\alpha = \langle C_0\rangle_\beta$ if $i = 0$. Then $\rho_j(\beta)$ is not $(t_{0,j}t_{1,j}\ldots t_{k-1,j})^\omega$ or its prefix. This is a contradiction. □

Suppose that a firing of $t_{i,j}$ represents the transfer of a "token" (or "privilege") from C_i to C_{i+1}. The condition that $\rho_j(\alpha)$ is either $(t_{0,j}t_{1,j}\ldots t_{k-1,j})^\omega$ or its prefix for any $\alpha \in \mathcal{L}(k)$ implies that there exists a unique token in R^k and initially the token resides in C_0. Theorems 1 and 2 state that if R^2 is structurally bounded, $R^2 \approx R^3$ and there exists a unique token in R^2, then there exists a unique token in R^k for any $k \geq 2$. We illustrate this proof method in Section 4.

4 Token-Passing Mutual Exclusion

Mutual exclusion is the problem of ensuring that at most one process among a set of k processes will be in its "critical section" at a time. One way to assure mutual exclusion is to let the processes form a ring and circulate a unique "privilege token" so that only the process that has the token can enter its critical section [13] [16]. Such a token-passing mutual

[3] We regard the reachability graph of a bounded Petri net with fairness as the state transition diagram of an ω-automaton that accepts both finite and infinite sequences [15], and then use known decision algorithms for such automata. Although the containment problem for ω-automata is PSPACE-complete [18] and thus the decision algorithms can be highly inefficient, it may still be feasible to use this method for small rings such as R^2 and R^3. Details will be reported elsewhere.

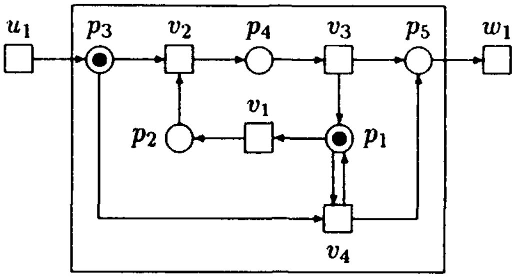

Figure 6: Component C for token-passing mutual exclusion.

Table 1: Places and transitions of C.

p_1	idle
p_2	waiting
p_3	have received the privilege token
p_4	critical section
p_5	ready to send the privilege token
u_1	receive the privilege token
v_1	request the critical section
v_2	enter the critical section
v_3	leave the critical section
v_4	pass the privilege token
w_1	send the privilege token

exclusion algorithm is used in [25] to illustrate the use of an invariant-based induction theorem. In this section, we model each process of a ring as a component and use our induction theorem (Theorem 1) to prove that the given algorithm is correct regardless of the size of the ring. We follow the general strategy outlined at the end of Section 3.

In this section, "C" refers to the component shown in Figure 6 that models a process in such a ring. Table 1 describes the events and conditions represented by the transitions and places. Transition u_1 is the only left interface transition, and w_1 the is only right interface transition. The initial marking of ring R^k is given as $M^k = (M, M', \ldots, M')$, where M is for C_0 and M' for $C_1, \ldots, C_{k-1}$. M is given by $M(p_1) = 1$, $M(p_2) = 0$, $M(p_3) = 1$, $M(p_4) = 0$ and $M(p_5) = 0$, which we write (10100). Using the same notation, we define $M' = (10000)$. Thus initially, all components are idling and C_0 has the unique privilege token in place $p_{0,3}$. ($p_{0,3}$ is the copy of p_3 in C_0.) We take $T^k = \emptyset$, and thus $f^k = f(T^k) = \mathbf{true}$. So a component can make either infinitely many requests or only a finite number of requests.

Component C fires v_1 when it requests the critical section and then waits (in p_2) until the privilege token arrives in place p_3 by a firing of u_1. Then it enters and leaves the critical section by firing v_2 and v_3, respectively. This brings the privilege token to p_5, and a firing of w_1 sends it to the next component. If the privilege token arrives in p_3 when C is idling, then it can be sent to p_5 by a firing of v_4. Note that progress assures that the privilege token eventually reaches p_5.

Lemma 8 *R^2 is structurally bounded.*

Proof The proof is found in Appendix A. □

The reachability graph G^k of R^k is a directed graph in which the vertices represent the markings of R^k reachable from the initial marking M^k and there is an arc with label t from a vertex v to vertex v' if the marking represented by v' is reachable from the marking represented by v when transition t fires. For convenience, we identify the markings of R^k and the vertices of G^k that represent them. (So "vertex M^k" refers to the vertex representing marking M^k.) Any firing sequence in R^k corresponds to a path in G^k in a natural way.

Lemma 9 $R^2 \approx R^3$.

Proof Since the set T^k of transitions that must be fired fairly is empty for all $k \geq 2$, we only need to show that $R^2 \sim R^3$. By Lemma 3, it suffices to show that

1. $\{[C_0]_\alpha \mid \alpha \in \mathcal{L}(2)\} = \{[C_0]_\alpha \mid \alpha \in \mathcal{L}(3)\}$, and
2. $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}(2)\} = \{[C_i]_\alpha \mid \alpha \in \mathcal{L}(3)\}$ for $i = 1, 2$.

In the following, we give an outline of the proof of $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}(2)\} = \{[C_1]_\alpha \mid \alpha \in \mathcal{L}(3)\}$ and leave the remaining cases to the reader. Since $T^2 = \emptyset$ and every vertex of G^2 (not shown) has at least one outgoing arc, $\mathcal{L}(2)$ consists of the infinite sequences represented by the infinite paths in G^2 starting from vertex M^2. By examining G^2, we can show that in any infinite path in G^2 starting from vertex M^2, arcs labeled $t_{0,1}$ and arcs labeled $t_{1,1}$ occur infinitely often and alternately, starting with an arc labeled $t_{0,1}$. ($t_{0,1}$ and $t_{1,1}$ are the copies of u_1 and w_1 in C_1, respectively.) So if we let $[[C_a \oplus \cdots \oplus C_b]]_\alpha$ denote the sequence obtained from $[C_a \oplus \cdots \oplus C_b]_\alpha$ by deleting the firability vectors, then we have $\{[[C_1]]_\alpha | \alpha \in \mathcal{L}(2)\} = \{(u_1 w_1)^\omega\}$. The firability vectors of C_1 have the form $\begin{bmatrix} x \\ y \end{bmatrix}$ where x and y are the token counts of places $p_{0,5}$ and $p_{1,5}$, respectively, and it is easy to insert them into $(u_1 w_1)^\omega$ to obtain

$$\{[C_1]_\alpha | \alpha \in \mathcal{L}(2)\} = \{(\begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1 \begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1)^\omega\}.$$

Using an analogous argument for R^3, we can show that

$$\{[C_1]_\alpha | \alpha \in \mathcal{L}(3)\} = \{(\begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \end{bmatrix} u_1 \begin{bmatrix} 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \end{bmatrix} w_1)^\omega\}.$$

Thus

$$\{[C_1]_\alpha | \alpha \in \mathcal{L}(2)\} = \{[C_1]_\alpha | \alpha \in \mathcal{L}(3)\}.$$

□

A firing sequence satisfies formula $t \supset \Diamond t'$ ("if t then eventually t'") if every occurrence of t is followed by an occurrence of t'.

Lemma 10 (Liveness of R^2) *A process that requests its critical section eventually enters it, i.e., for $i = 0, 1$,*

$$S_i \subseteq \{\langle\langle C_i \rangle\rangle_\alpha | \alpha \in \mathcal{L}(2)\} \subseteq S'$$

for some nonempty S_i and $S' = \{\alpha \in T^\omega | \alpha \text{ satisfies } v_1 \supset \Diamond v_2\}$, where T is the set of transitions of C.

Proof We prove the claim for $i = 1$ and leave the case $i = 0$ to the reader. The first "$\subseteq$" is trivial. We can show that every maximal simple path in G^2 starting with an arc labeled $v_{1,1}$ contains an arc labeled $v_{1,2}$. ($v_{1,1}$ and $v_{1,2}$ are the copies of v_1 and v_2 in C_1, respectively.) This proves the second "$\subseteq$." □

In R^2, a firing of $t_{0,1}$ (the copy of u_1 in C_1) and a firing of $t_{1,1}$ (the copy of w_1 in C_1) represent the transfer of the privilege token, and C_0 and C_1 can be in its critical section only while it has the privilege token. The following lemma is based on this observation. For $\alpha \in \mathcal{L}(2)$, $\rho(\alpha)$ denotes the sequence obtained from α by deleting all transitions except the interface transitions $t_{0,1}$ and $t_{1,1}$.

Lemma 11 (Safeness of R^2) *C_0 and C_1 cannot be in their critical sections at the same time, i.e., $\rho(\alpha) = (t_{0,1}t_{1,1})^\omega$ for any $\alpha \in \mathcal{L}(2)$.*

Proof The lemma is immediate from $\{[[C_1]]_\alpha | \alpha \in \mathcal{L}(2)\} = \{(u_1 w_1)^\omega\}$ given in the proof of Lemma 9. □

Finally, we have the following theorem.

Theorem 3 (Correctness of R^k) *For any $k \geq 2$, in ring R^k*

1. *a process that requests its critical section eventually enters it, and*

2. *no two processes can be in their critical sections at the same time.*

Proof The theorem follows from Theorems 1, 2 and Lemmas 8, 9, 10 and 11. □

5 A Simple Producer-Consumers System

Consider a ring consisting of one "producer" and many identical "consumers." The producer generates a product that is circulated in the ring. A consumer receiving a product can either pass it (without consuming it) to its right neighbor, or "consume" it and send "garbage" to the right neighbor. Garbage received by a consumer is always passed to to its right neighbor. The producer can generate a new product only when it receives garbage from its left neighbor. We assume that the producer is allowed to pass or consume a product that has been returned. If the producer consumes a product, it then sends garbage to its right neighbor. We assume that at any time, there can be only one object (a product or garbage) in the ring.

In this section, "C" refers to the component shown in Figure 7 that models a process in such a ring. We assume that in R^k consisting of k components, C_0 is the producer and $C_1, \ldots, C_{k-1}$ are the consumers. Table 2 describes the events and conditions represented

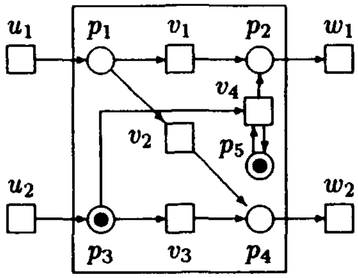

Figure 7: Component C for a producer-consumers system.

Table 2: Places and transitions of C.

p_1	have received a product
p_2	ready to send a product
p_3	have received garbage
p_4	ready to send garbage
p_5	one token for producer, empty tor consumer
u_1	receive a product
u_2	receive garbage
v_1	pass a product
v_2	consume a product
v_3	pass garbage
v_4	generate a product (producer only)
w_1	send a product
w_2	send garbage

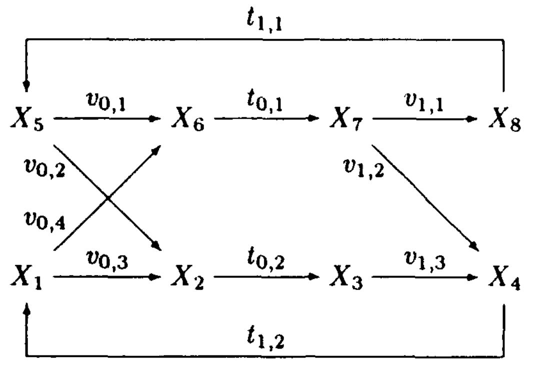

Figure 8: Structure of G^2.

by the transitions and places. Transition u_1 and u_2 are the left interface transitions, and w_1 and w_2 are the right interface transitions. The initial marking of ring R^k is given as $M^k = (M, M', \ldots, M')$, where $M = (00101)$ is for producer C_0 and $M' = (00000)$ is for consumers $C_1, \ldots, C_{k-1}$. (This notation was introduced in Section 4.) Note that transition v_4 ("generate a product") can fire only in C_0, and initially C_0 has garbage in place $p_{0,3}$. ($p_{0,3}$ is the copy of p_3 in C_0.) We take T^k to be the set of all transitions of R^k. This means that no component is allowed to always pass or always consume a product from some time on, and the producer must generate a product infinitely often if garbage is returned infinitely often. The system is considered to be correct if all components consume a product infinitely many times.

Lemma 12 *R^2 is structurally bounded.*

Proof The proof is found in Appendix A. □

Lemma 13 $R^2 \approx R^3$.

Proof By Lemma 3, it suffices to show that

1. $\{[C_0]_\alpha \mid \alpha \in \mathcal{L}_{\neg 0}(2)\} = \{[C_0]_\alpha \mid \alpha \in \mathcal{L}_{\neg 0}(3)\}$, and
2. $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\} = \{[C_i]_\alpha \mid \alpha \in \mathcal{L}_{\neg i}(3)\}$ for $i = 1, 2$.

In the following, we give an outline of the proof of

$$\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\} = \{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(3)\} \tag{5}$$

and leave the remaining cases to the reader. As we did in the proof of Lemma 9, let us first characterize the set $\{[[C_1]]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\}$. Figure 8 shows the structure of G^2 and the labels of its arcs, where vertex X_1 represents the initial marking M^2. Since every vertex of G^2 has at least one outgoing arc, no finite path in G^2 represents a firing sequence in $\mathcal{L}_{\neg 1}(2)$. This, together with the structure of G^2 and the fact that $t_{0,1}$, $t_{0,2}$, $t_{1,1}$ and $t_{1,2}$ are the copies of u_1, u_2, w_1 and w_2 in C_1, respectively, shows that $\{[[C_1]]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\}$

is a subset of $\{u_1w_1, u_1w_2, u_2w_2\}^\omega$. Furthermore, any infinite path from X_1 representing a firing sequence in $\mathcal{L}_{\neg 1}(2)$ must visit vertex X_1 infinitely often (by the firings of $t_{1,2}$), since otherwise, X_5 is visited infinitely often but X_2 is not, and thus the fairness condition at C_0 (that $v_{0,2}$ must be fired infinitely often if it becomes firable infinitely often) is violated. Since X_1 is visited infinitely often, again by the fairness condition on $v_{0,3}$ and $v_{0,4}$ of C_0, both $t_{0,1}$ and $t_{0,2}$ must fire infinitely often. Also, the fairness condition on $v_{0,1}$ and $v_{0,2}$ of C_0 requires that if X_5 is visited infinitely often (by the firings of $t_{1,1}$), then both $v_{0,1}$ and $v_{0,2}$ must fire infinitely often. Thus $\{[[C_1]]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\} \subseteq U$, where U is the set of sequences in $\{u_1w_1, u_1w_2, u_2w_2\}^\omega$ such that (a) both w_2u_1 and w_2u_2 appear infinitely often, and (b) if w_1 appears infinitely often then both w_1u_1 and w_1u_2 appear infinitely often. Conversely, we can easily show that for any sequence $\sigma \in U$, there exists some $\tau \in \mathcal{L}_{\neg 1}(2)$ that is legal at C_0 such that $[[C_1]]_\tau = \sigma$. Therefore $\{[[C_1]]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\} = U$. We then obtain $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\}$ by inserting, into the sequences in U, the firability vectors of C_1 having the form

$$\begin{bmatrix} x_1 \\ x_2 \\ x_3 \\ x_4 \end{bmatrix}$$

where $x_1, \ldots, x_4$ are the token counts of places $p_{0,2}$, $p_{0,4}$, $p_{1,2}$ and $p_{1,4}$, respectively. Since (a) at most one of $p_{0,2}$, $p_{0,4}$, $p_{1,2}$ and $p_{1,4}$ can have a token at a time and (b) $p_{0,2}$, $p_{0,4}$, $p_{1,2}$ and $p_{1,4}$ can lose a token only when $t_{0,1}$, $t_{0,2}$, $t_{1,1}$ and $t_{1,2}$ fire, respectively, $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\}$ is the set of sequences obtained from the sequences in U by replacing u_1, u_2, w_1 and w_2 by $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} 1 \\ 0 \\ 0 \\ 0 \end{bmatrix} u_1$, $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 1 \\ 0 \\ 0 \end{bmatrix} u_2$, $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 0 \\ 1 \\ 0 \end{bmatrix} w_1$ and $\begin{bmatrix} 0 \\ 0 \\ 0 \\ 0 \end{bmatrix}\begin{bmatrix} 0 \\ 0 \\ 0 \\ 1 \end{bmatrix} w_2$, respectively. Using an analogous argument on G^3 that has 12 vertices, we can show that $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(3)\}$ coincides with $\{[C_1]_\alpha \mid \alpha \in \mathcal{L}_{\neg 1}(2)\}$ obtained above. □

A firing sequence satisfies formula $\Box\Diamond t$ ("infinitely often t") if t occurs infinitely often in it.

Lemma 14 (Liveness of R^2) *Both C_0 and C_1 consume a product infinitely often, i.e., for $i = 0, 1$,*

$$S_i \subseteq \{\langle\langle C_i \rangle\rangle_\alpha | \alpha \in \mathcal{L}(2)\} \subseteq S'$$

for some nonempty S_i and $S' = \{\alpha \in T^\omega | \alpha$ satisfies $\Box\Diamond v_2\}$, where T is the set of transitions of C.

Proof By examining G^2 and using the fairness condition on C_0 and C_1, we can show that both $v_{0,2}$ and $v_{1,2}$ must fire infinitely often, where $v_{0,2}$ and $v_{1,2}$ are the copies of v_2 ("consume a product") in C_0 and C_1, respectively. The argument is basically similar to that used in the proof of Lemma 13, and is thus omitted. □

Finally, we have the following theorem.

Theorem 4 (Correctness of R^k) *For any $k \geq 2$, in ring R^k each component consumes a product infinitely often.*

Proof The theorem follows from Theorem 1 and Lemmas 12, 13 and 14. □

6 Concluding Reinarks

We have introduced the concept of similarity between two process rings of fair Petri nets, and proved a structural indiction theorem (Theorem 1) that can be used to prove the correctness of a ring of any large size from the correctness of a ring having only a few components. The theorem has been applied to the verification problem of two examples, token-passing mutual exclusion and a simple producer-consumers system.

The main condition needed for applying the theorem is the strong similarity between R^2 and R^3, i.e., $R^2 \approx R^3$. It can happen, however, that for some $k \geq 3$, all rings R^ℓ, $\ell \geq k$, are mutually similar but $R^2 \approx R^3$ does not hold. Such rings may still admit induction similar to that of Theorem 1. We are currently working on a more general version of the theorem that can be applied to such cases. Some results in this direction can be found in [9].

As was pointed out in Section 3, testing the strong similarity of two rings using an automatic verifier can be time consuming. It is desirable that we find simple sufficient conditions for two rings to be strongly similar. Another direction of research is to apply the ideas developed for rings in this paper to other network topologies, such as stars, trees, chains, meshes and completely connected graphs. It is an interesting problem to develop analogous induction methods for such networks.

Appendix A

For a Petri net $N = (P, T, F)$ such that $P = \{p_1, \ldots, p_n\}$ and $T = \{t_1, \ldots, t_m\}$, the *incidence matrix* of N is an $m \times n$ matrix $A = [a_{i,j}]$ such that $a_{i,j} = F(t_i, p_j) - F(p_j, t_i)$. Note that $a_{i,j}$ is the change in the token count of place p_j when transition t_i fires once. It is known that N is structurally bounded iff there exists an n-dimensional vector y of positive integers such that $Ay \leq 0$ [14]. (y is called an *S-invariant* if $Ay = 0$.) The condition $Ay \leq 0$ assures that the weighted sum of token counts of a marking never increases after a firing of any transition, where the j-th element of y is the weight assigned to p_j.

Proof of Lemma 2 Assume that C has n places, s internal transitions and m interface transitions on each side. For any $k \geq 2$, since the components $C_0, \ldots, C_{k-1}$ have the same structure and only the interface transitions between two components can be connected to the places in both, the incidence matrix A_k for R^k can be written as a $k(s+m) \times kn$ matrix

$$A_k = \begin{bmatrix} B & \mathbf{0} & \mathbf{0} & \cdots & \mathbf{0} \\ D & E & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & B & \mathbf{0} & \cdots & \mathbf{0} \\ \mathbf{0} & D & E & \cdots & \mathbf{0} \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \cdots & B \\ E & \mathbf{0} & \mathbf{0} & \cdots & D \end{bmatrix}$$

where B is an $s \times n$ matrix describing the connections among the n places and s internal transitions of a component, D and E are $m \times n$ matrices such that $(D\ E)$ describes the connections among the $2n$ places of two consecutive components and m interface transitions between them, and $\mathbf{0}$ is a zero matrix of appropriate dimensions. Since R^2 is structurally bounded, there exists a $2n$-dimensional vector of positive integers

$$y_2 = \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix}$$

where $\mathbf{a}$ and $\mathbf{b}$ are n-dimensional vectors, such that

$$A_2 y_2 = \begin{bmatrix} B & \mathbf{0} \\ D & E \\ \mathbf{0} & B \\ E & D \end{bmatrix} \begin{bmatrix} \mathbf{a} \\ \mathbf{b} \end{bmatrix} \leq 0.$$

Then $B\mathbf{a} \leq 0$, $D\mathbf{a} + E\mathbf{b} \leq 0$, $B\mathbf{b} \leq 0$, and $E\mathbf{a} + D\mathbf{b} \leq 0$, and thus

$$A_2 \begin{bmatrix} \mathbf{a}+\mathbf{b} \\ \mathbf{a}+\mathbf{b} \end{bmatrix} \leq 0.$$

Then it is easy to show that the kn-dimensional vector

$$y_k = \begin{bmatrix} \mathbf{a}+\mathbf{b} \\ \mathbf{a}+\mathbf{b} \\ \vdots \\ \mathbf{a}+\mathbf{b} \end{bmatrix}$$

satisfies

$$A_k y_k \leq 0.$$

So R^k is structurally bounded. □

Proof of Lemma 8 Choose y that assigns 2 to $p_{0,4}$ and $p_{1,4}$ (the copies of p_4 in C_0 and C_1, respectively), and 1 to all other places. Then y satisfies $A_2 y = 0$, where A_2 is the incidence matrix of R^2. □

Proof of Lemma 12 Choose y that assigns 1 to all places. Then y satisfies $A_2 y \doteq 0$, where A_2 is the incidence matrix of R^2. □

Acknowledgment We wish to thank the anonymous referee for careful reading of the manuscript and many helpful comments.

References

[1] K. R. Apt and D. C. Kozen, "Limits for automatic verification of finite-state concurrent systems," *Information Processing Letters 15*, May 1986, pp. 307-309.

[2] L. A. Cherkasova and V. E. Kotov, "The undecidability of propositional temporal logic for Petri nets," *Computers and Artificial Intelligence 6*, No. 2, 1987.

[3] E. M. Clarke, O. Grümberg and M. C. Browne, "Reasoning about networks with many identical finite-state processes," *Proceedings of the 5th Annual ACM Symposium on Principles of Distributed Computing*, Calgary, Alberta, Canada, August 1986, pp. 240–248.

[4] E. W. Dijkstra, "Hierarchical ordering of sequential processes," *Acta Informatica 1*, 2, 1971, pp. 115–138.

[5] S. German and A.P. Sistla, "Reasoning about systems with many processes," Te‹hnical Report TR88-378.02, GTE Laboratories, March 1988.

[6] R. R. Howell and L. E. Rosier, "Problems concerning fairness and temporal logic for conflict-free Petri nets," *Theoretical Computer Science 64*, 1989, pp. 305–329.

[7] R. R. Howell, L. E. Rosier and H. C. Yen, "A taxonomy of fairness and temporal logic problem for Petri nets," *Theoretical Computer Science 82*, 1991, pp. 341–372.

[8] R. P. Kurshan and K. McMillan, "A structural induction theorem for processes," *Proc. 8th ACM Symp. Principles of Distributed Computing*, Edmonton, Alberta, August 1989, pp. 239–247.

[9] J. Li, I. Suzuki and M. Yamashita, "A new structural induction theorem for rings of temporal Petri nets," to appear in *IEEE Transactions on Software Engineering.*

[10] H. Lu and I. Suzuki, "Application of temporal Petri nets to verification of handshake daisy chain arbiters," in *Proceedings of the 29th Midwest Symposium on Circuits and Systems*, Lincoln, Nebraska, 1987, pp. 823–826.

[11] Z. Manna and A. Pnueli, "Verification of concurrent programs: the temporal framework," in *The Correctness Problem in Computer Science*, R. S. Boyer and J. S. Moore eds., International Lecture Series in Computer Science, Academic Press, New York, 1981, pp. 215–273.

[12] Z. Manna and P. Wolper, "Synthesis of communi^cation processes from temporal logic specifications," *ACM Tr^ansactions on Programming Languages and Systems 6* 2, January 1984, pp. 68–93.

[13] A. J. Martin, "Distributed mutual exclusion on a ring of processes," *Science of Computer Programming 5*, 1985, pp. 265–276.

[14] T. Murata, "Petri nets: Properties, analysis and applications," *Proceedings of the IEEE 77*, No. 4, 1989, pp. 541–580.

[15] D. Park, "Concurrency and automata on infinite sequences," in *Proceedings of the 5th GI-Conference on Theoretical Computer Science*, Lecture Notes in Computer Science 104, Springer Verlag, 1981, pp. 167–183.

[16] W. W. Plummer, "Asynchronous arbiters," *IEEE Transactions on Computers C-21*, January 1972, pp. 37–42.

[17] A. Pnueli, "The temporal logic of programs," *Proc. 18th Symp. on Foundations of Computer Science*, Province, November 1977, pp. 46–57.

[18] A. P. Sistla, M. Y. Vardi and P. Wolper, "The complementation problem for Buchi automata with applications to temporal logic," *Theoretical Computer Science 49*, 1987, pp. 217–237.

[19] I. Suzuki, "Fundamental properties and applications of temporal Petri nets," in *Proceedings of the 9th Annual Conference on Information Sciences and Systems*, Johns Hopkins University, Baltimore, Maryland, 1985, pp. 641–646.

[20] I. Suzuki, "Proving properties of a ring of finite state machines," *Information Processing Letters 28*, 1988, pp. 213–214.

[21] I. Suzuki, "Formal analysis of the alternating bit protocol by temporal Petri nets," *IEEE Transactions on Software Engineering 16*, No. 11, November 1990, pp. 1273–1281.

[22] I. Suzuki and H. Lu, "Temporal Petri nets and their application to modeling and analysis of a handshake daisy chain arbiter," *IEEE Transactions on Computers 38*, No. 5, 1989, pp. 696–704.

[23] N. Uchihira and S. Honiden, "Verification and synthesis of concurrent programs using Petri nets and temporal logic," *The Transactions of the IEICE E73*, No. 12, 1990, pp. 2001–2010.

[24] R. Valk, "Infinite behavior of Petri nets," *Theoretical Computer Science 25*, 1983, pp. 311-341.

[25] P. Wolper and V. Lovinfosse, "Verifying properties of large sets of processes with network invariants," in *Automatic Verification Methods for Finite State Systems*, J. Sifakis (ed.), Lecture Notes in Computer Science 407, Springer Verlag, 1990, pp. 68–80.